

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
28 June 2001 (28.06.2001)

PCT

(10) International Publication Number
WO 01/47159 A1

(51) International Patent Classification⁷: H04J 3/04 (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(21) International Application Number: PCT/US00/34173

(22) International Filing Date: 15 December 2000 (15.12.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 09/468,192 21 December 1999 (21.12.1999) US

(71) Applicant: SYCAMORE NETWORKS, INC. [US/US]; 150 Apollo Drive, Chelmsford, MA 01824-0986 (US).

(72) Inventors: AZIZOGLU, Murat; 5 Amherst Street, North Billerica, MA 01862 (US). KONG, Thomas; 319 St. Paul Street, Apt. 1, Brookline, MA 02446 (US). SWANSON, Eric, A.; 13R Oakwood Road, Acton, MA 01721 (US).

(74) Agents: SCHURGIN, Stanley, M. et al.; Weingarten, Schurgin, Gagnebin & Hayes, LLP, Ten Post Office Square, Boston, MA 02109 (US).

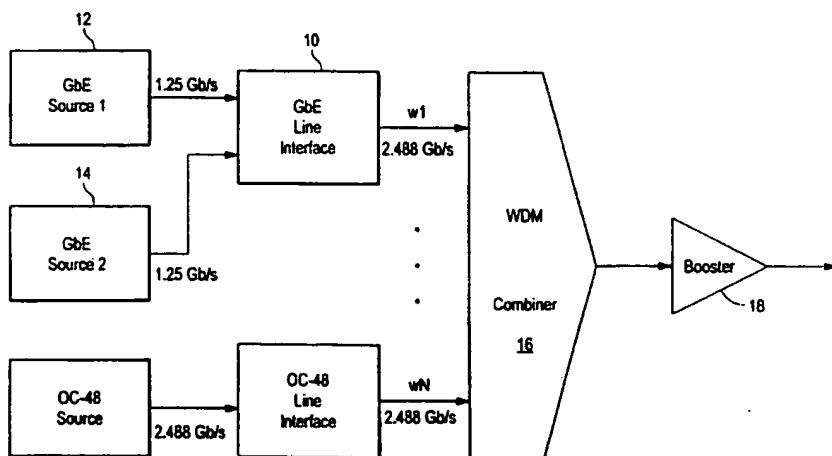
(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— With international search report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR TRANSPORTING GIGABIT ETHERNET AND FIBER CHANNEL SIGNALS IN WAVELENGTH-DIVISION MULTIPLEXED SYSTEMS



WO 01/47159 A1

(57) Abstract: Multiple gigabit Ethernet (GbE) and Fiber Channel (FC) signals are multiplexed (16) and transported on a wavelength-division multiplexed (WDM) communications link using synchronous optical network (SONET) signaling. A transmitter decodes the 8b/10b encoded GbE/FC signals to reduce their respective signaling rates to no greater than the payload data rate of an OC-48 signal used on the link. The decoded signals are packetized into fixed-size packets in order to append a stream identifier used for de-multiplexing at the receive end. The packetized streams are multiplexed (16) together and formatted into the OC-48 signal, which is transmitted over the communications link. A receiver de-multiplexes the received signal to obtain the packet streams, converts the packet streams to the corresponding decoded signals, and re-encodes each signal to re-create the original GbE/FC signals. A pair of interfaces can be arranged to provide independent adding or dropping of the GbE/FC signals that share the same wavelength.

-1-

TITLE OF THE INVENTION

**Method And Apparatus For Transporting Gigabit Ethernet
And Fiber Channel Signals In Wavelength-Division
Multiplexed Systems**

5

CROSS REFERENCE TO RELATED APPLICATIONS

--Not Applicable--

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR
10 DEVELOPMENT

--Not Applicable--

BACKGROUND OF THE INVENTION

15 The present invention relates to the transport of packet traffic in optical communications networks employing synchronous signaling techniques, such as networks employing the Synchronous Optical Network (SONET) signaling format.

20 There exist today a large number of wavelength-division multiplexed (WDM) point-to-point transmission systems and WDM networks. Many of these systems have been designed to support a fixed number of wavelengths and a predetermined data format. SONET is a common framing format for data transport in long-haul and
25 metropolitan-area WDM carrier networks. The SONET frame provides a convenient standard mechanism to multiplex and transport circuit-switched traffic in high-speed backbones. There are a number of readily-available standard components for SONET-based systems, such as

-2-

clock recovery units designed for SONET rates, SONET framers and multiplexers, so that the task of network design is simplified. SONET also provides mechanisms to support network functions such as error detection, alarm 5 insertion, automatic protection switching, etc. The common frame format also allows for advanced functions, such as electro-optic switching, signal regeneration, internal signaling, and signal restoration, that are independent of the format of the native traffic carried 10 by the network. Because of these and other beneficial characteristics, SONET framing is used in most WDM systems deployed today.

In contrast, many packet-switched local area networks (LANs) use framing defined in the 15 long-established Ethernet standard. Unlike SONET, Ethernet and other LAN protocols rely on non-synchronous signaling techniques. While the original Ethernet standard contemplated the use of copper coaxial cable as a physical transmission medium, there have been recent 20 extensions, such as Gigabit Ethernet (GbE), that contemplate the use of optical media. GbE is an evolution of the Ethernet LAN standard to gigabit rates. It uses the same frame format specified by the original Ethernet standard, and the same multiple-access protocol 25 and flow-control methods as the 10 Mb/s and 100 Mb/s Ethernet standards. GbE also employs the same variable frame length (64-1514 byte packets) specified in the Ethernet and Fast Ethernet standards. This backward compatibility makes it easier to connect existing lower- 30 speed Ethernet devices to GbE devices using LAN switches and routers for speed adaptation.

-3-

The GbE standard defines both the physical (PHY) and the medium access control (MAC) layers of the OSI architecture. The PHY layer deals with the transmission of bits over physical channels. In particular, the GbE 5 standard specifies a raw bit rate of 1.25 Gb/s. In full-duplex mode, GbE employs a packet switch at the network hub that operates without any collisions. While operating in half-duplex mode, GbE uses the classical CSMA/CD mechanism to resolve contentions for access to 10 the shared physical transmission medium, although the slot duration has been extended over that of Ethernet and Fast Ethernet to allow transmitters sufficient time for collision detection.

Many of the PHY layer features of GbE also appear 15 in another standard called Fiber Channel (FC), which is a protocol for point-to-point high-speed computer interconnection. The FC protocol precedes GbE by several years and has become very popular in practice. FC defines a PHY layer protocol for the point-to-point 20 interconnection of two nodes at a rate of 1.0625 Gb/s.

The FC standard specifies a line encoding algorithm known as "8b/10b" encoding to achieve DC balance and "run length limiting", i.e., providing a minimum rate 25 of signaling transitions in the data stream to ensure adequate clock recovery at a receiver. The 8b/10b code converts a byte-wide data stream of random '0's and '1's into a DC-balanced binary stream with a maximum run length of five '1's or five '0's. This is accomplished by converting 8-bit blocks of the non-encoded data 30 stream into 10-bit code words selected to achieve the run-length limitation. The encoding and decoding algorithms can be readily implemented. At the encoder,

-4-

two different sub-codes operate on 5-bit and 3-bit sub-blocks of each input block to yield corresponding 6-bit and 4-bit sub-codewords, with minimal interaction between the 5b/6b and the 3b/4b sub-codewords. The code 5 constrains the disparity between the number of 1s and 0s in a sub-block to be -2, 0, or +2. Furthermore, by assigning two sequences to certain inputs, and by switching between these two sequences according to the state of the encoder, the encoder limits the running 10 disparity observed at sub-block boundaries to +1, 0, or -1. Thus, the 8b/10b code thus provides a mechanism for efficient clock recovery and DC balance, albeit at the expense of a 25% increase in the raw signalling rate of the channel.

15 Because SONET and GbE have been separately optimized for transport and data networking, respectively, the existing art has treated these signaling mechanisms in an isolated manner. Thus, a typical WDM network with OC-48 interfaces (a particular 20 type of SONET signal) can accept only SONET-framed traffic at the OC-48 rate (2.488 Gb/s); there is no ready mechanism for accepting variable-sized packets transmitted at a non-SONET rate, such as the 1.25Gb/s rate of a GbE network. While it is possible to use a 25 device known as a "SONET framer" to translate a GbE signal into an OC-48 SONET signal, this approach would result in poor utilization (about 40%) of the available bandwidth at each wavelength transporting such traffic.

30 An alternative optical transport strategy is to use a so-called "transparent interface", which foregoes SONET framing and performs a simple frequency translation at the optical network boundary. The

-5-

advantage of this approach is bit rate and format independence. However, this approach suffers from drawbacks such as the lack of performance monitoring within the optical network; the inability to perform 5 opto-electronic circuit switching; and increased jitter accumulation in regenerator cascades.

There is a need for an optical network interface that can accept and multiplex multiple GbE/FC signals into a synchronous format signal such as a SONET signal, 10 in order to provide for multiplexing, adding/dropping, and monitoring LAN traffic in an optical backbone.

BRIEF SUMMARY OF THE INVENTION

In accordance with the present invention, methods 15 and apparatus are disclosed to perform multiplexing and transport of Gigabit Ethernet (GbE) and Fiber Channel (FC) signals on WDM networks using SONET signaling. A line interface in an optical communications node is described that accepts multiple GbE/FC streams, 20 processes these individual streams to reduce their data rates, and combines the rate-matched stream into a standard-rate SONET stream which will be transported on a wave. The data-carrying capacity of the SONET network is efficiently utilized, and important functionality 25 such as synchronization and performance monitoring available in SONET systems is provided.

A transmitter decodes encoded data signals such as GbE or FC signals to generate a corresponding plurality of non-encoded data signals, so that the sum of the 30 respective signaling rates of the non-encoded data signals is no greater than the data-carrying capacity of the communications link. In one embodiment this

-6-

decoding reduces the signaling rate of each signal to about 9/10 the signaling rate of the encoded signal. The non-encoded data signals are multiplexed together to form a multiplexed data signal, the multiplexed data 5 signal is transmitted over the communications link. In one embodiment the multiplexed signal is an OC-48 or similar synchronous transport signal. A receiver receives the multiplexed data signal from the communications link, de-multiplexes the received signal 10 to obtain the non-encoded data signals, and re-encodes each de-multiplexed data signal to re-create the encoded data signals at the receive end.

To enable correct de-multiplexing at the receive end, the disclosed system employs packetization of the 15 signals after decoding, where each packet includes an address or stream identifier. The transmitter converts sequential data blocks of each non-encoded signal into corresponding packets, and asynchronously interleaves the packets of the non-encoded signals to create the 20 multiplexed data signal. The packets may be of a fixed size, which is chosen to achieve a desired balance between link bandwidth efficiency and storage efficiency. In one embodiment a 67 byte packet size is used. The receiver uses the stream identifier of the 25 packets to separate the streams for the different signals, and then re-encodes each signal to re-create the encoded signals at the receive end.

The stripping of coding overhead at a transmitter and re-encoding at the receiver advantageously permits 30 multiple signals, such as multiple GbE or FC signals, to be carried on a single transport signal such as an OC-48 signal. All the advantages of SONET transport, such as

-7-

error monitoring, flexible add/drop, etc., can be realized for the GbE or FC traffic. Moreover, because SONET includes a mechanism for maintaining synchronization, accurate end-to-end operation is obtained despite the removal of the timing information embedded from the encoded data signals before transport over the communications link.

Other aspects, features, and advantages of the present invention are disclosed in the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

Figure 1 is a block diagram of the transmit side of a wavelength-division multiplexed (WDM) optical communications node in accordance with the present invention;

Figure 2 is a high-level block diagram of the transmit side of a high-speed line interface in the WDM optical communications node of Figure 1;

Figure 3 is a more detailed block diagram of the transmit side of the line interface shown in Figure 2;

Figure 4 is a detailed block diagram of the receive side of a high-speed line interface in the WDM optical communications node of Figure 1;

Figure 5 is a block diagram showing an interconnection among two line interfaces such as shown in Figures 3 and 4;

Figure 6 is a block diagram of the transmit side of a high-speed line interface similar to that of Figure 3 using forward-error-correction (FEC) encoding; and

-8-

Figure 7 is a block diagram of the receive side of a high-speed line interface similar to that of Figure 4 using forward-error-correction (FEC) encoding.

5

DETAILED DESCRIPTION OF THE INVENTION

Figure 1 shows an example of a transmitter node in a wavelength-division multiplexed (WDM) system. The WDM node includes a GbE line interface 10 that multiplexes two GbE signals received from respective GbE sources 12, 14 onto a single optical signal having wavelength w1. As shown, signals from other sources are mapped into other wavelengths w2...wN. The line interface 10 converts the input signals into a SONET signal (OC-48 in this example), which is used to modulate an optical carrier of wavelength w1, which in the illustrated embodiment is in the range of 1530-1605 nm. The different wavelengths are combined using a WDM combiner 16. The output of the combiner 16 is sent to a booster amplifier 18, and the amplified WDM signal is supplied to a fiber (not shown) for transmission to a remote node. The transmitted signal typically travels through a series of transmission spans, each terminated by an optical amplifier.

25

There is a fundamental mismatch between the 1.25 Gb/s signaling rate of a GbE signal and the 2.488 Gb/s capacity of an OC-48-formatted wave. If an OC-48 wave were populated by only a single GbE, slightly less than half the wave capacity is wasted. However, an OC-48 wave has insufficient capacity to carry two such GbE signals. In the illustrated system, however, a technique is used that enables one OC-48 wave to carry two GbE or FC signals.

-9-

A multiplexing method for combining two GbE signals into an OC-48 wavelength is shown in Figure 2. Each of the two GbE inputs is first converted into a 10-bit parallel stream by a corresponding 5 serializer/deserializer (Ser/Des) 20-1 or 20-2. The 10-bit parallel streams are then decoded by 8b/10b codecs 22-1 and 22-2, which remove the run-length code overhead from each stream. Each codec 22 provides error 10 information to performance monitoring (PM) logic 24 which monitors the bit error rate (BER) performance of the GbE signals. The output of each codec 22 is an 8-bit 15 parallel stream at a bit rate of 1 Gb/s. A ninth bit is added to indicate whether a byte has originated from a special type of character called a "K-character" used in GbE systems. GbE employs these K-characters for the exchange of control information, such as the indication of start and end of a packet, idle channel status, etc. Since all 256 8-bit characters occur within the data, a ninth bit is necessary to indicate the presence of a 20 K-character. As a result, the aggregate rate at the output of the two 8b/10b codecs 22 is 2.25 Gb/s. The rate-reduced streams are supplied to multiplexing and framing logic 26, which multiplexes the streams together using an asynchronous statistical multiplexing technique 25 described below, and maps the multiplexed streams into synchronous frames in the OC-48 signal.

An implementation of the transmitter is shown in more detail in Figure 3. An interleaver 30 is shown that can be implemented using a field programmable gate array (FPGA). The data bytes from the 8b/10b codecs 22 are 30 stored in FIFO buffers 32 in the interleaver FPGA 30. Interleaving logic 34-1 packetizes the data streams from

-10-

the FIFOs 32-1 and 32-2 into fixed-size packets, and interleaves the streams using statistical multiplexing. This function is also performed by interleaving logic 34-2 for separate data streams, as described below. If 5 both FIFO buffers 32-1 and 32-2 are non-empty, packets are selected in an alternating manner from the FIFOs 32-1 and 32-2 and sent successively to a SONET framer device 36 via multiplexers 38 and 40. If one of the FIFOs 32-1 or 32-2 is empty while the other is 10 non-empty, then the interleaving logic 34-1 pulls multiple packets from the non-empty FIFO for successive insertion into the OC-48 stream, thus making full use of available transmit bandwidth. As shown, the output of the SONET framer 36 is a 16-bit parallel signal, which 15 is converted by a serializer 42 into a serial OC-48 signal used to modulate a laser 44.

The statistical time division multiplexing (TDM) technique requires the use of a stream identifier in each transmitted packet to enable proper de-multiplexing 20 at the receiving end. A framing byte (or multiple bytes) are also needed. Both the stream identifier and framing bytes represent overhead that consumes bandwidth. The size of the packet is selected to achieve a suitable compromise between bandwidth 25 efficiency (which tends to increase with greater packet size for a given amount of overhead) and storage efficiency (which tends to decrease with packet size for a given average GbE frame size). Bandwidth efficiency is important when there is only a small margin between 30 the aggregate data rate of the signals being combined (e.g., two GbE signals) and the payload data rate of the carrier signal (e.g., the OC-48 signal). Storage

-11-

efficiency is important in applications having bursty traffic patterns, such as a GbE LAN. If the packets are too large, latency in the FIFO buffers 32 is undesirably large. In the illustrated embodiment, a 67-byte packet
5 is used to achieve the desired balance between bandwidth and storage efficiency. The packet has a 64-byte payload, a 1-byte address or stream identifier, a 1-byte header flag and a 1-byte trailer flag. The resulting bandwidth expansion keeps the aggregate bit rate below
10 that of the OC-48 payload rate while minimizing the storage required in the FPGA 30.

Figure 3 also shows how line-side protection capability can be supported by the GbE line interface 10, either in a "1+1" mode or in a "1:1" mode. At the
15 transmitting node, the two GbE signals are sent via an electrical interface 46 to a neighboring line interface card (not shown) which operates at the same wavelength, but in the opposite direction on the ring. Similarly, the neighboring card forwards its two GbE signals to the line interface card 10. Thus, the transmitter side of
20 the line interface card 10 has four GbE inputs: inputs 1 and 2 via respective 1310 nm optical interfaces (or any other suitable physical interface), and inputs 3 and 4 via respective electrical connections from the neighboring line interface card. In 1+1 protection mode,
25 either the input pair (1,2) are active or the input pair (3,4) are active, depending on whether the line interface card 10 is part of the working path or the protection path. The active pair is selected by the multiplexer 38. In 1:1 protection mode, all four input ports are active. Under normal operation, GbE signals from input ports 3 and 4 are simply terminated at the
30

-12-

interleaver FPGA 30. When a failure occurs, the multiplexer 38 selects the signals from ports 3 and 4 and terminates the signals from ports 1 and 2.

5 As shown in Figure 3, the interleaver FPGA also contains an idle generator 48 and a bit error rate tester (BERT) 50 that are used for testing and monitoring purposes. The multiplexer 40 is used to select from among the normal data stream or the respective outputs of these devices.

10 Figure 4 shows a functional block diagram of the receive side of the GbE line interface 10. The output from the avalanche photo-diode (APD) 60 is supplied to a de-serializer 62. The output of the de-serializer 62 is provided to the SONET framer 36, which removes SONET 15 overhead and carries out performance monitoring functions. A de-interleaver FPGA 64 separates packets belonging to different streams using the address byte in the packet headers. After the appropriate stream from the working or protection side is chosen, the 8b/10b 20 codecs 22 re-insert the run-length code. Thus, from an input/output perspective, the removal of the run-length coding for WDM transport is invisible to the GbE or FC devices that send and receive the GbE/FC signal carried over the WDM optical link.

25 The method described here is also applicable when one of the inputs is a GbE signal and the other is a FC signal. Since the two inputs operate on separate clocks, the appropriate clock rates (one at 1.25 GHz and one at 1.0625 GHz) enable the operation with hybrid GbE/FC 30 inputs without any modifications to the transmitter/receiver hardware. It will be appreciated that between a neighboring pair of line

interface cards, four pairs of signals are routed: two pair of GbE/FC signals (ports 1,2 and ports 3,4) for the transmit side, and two pair of GbE/FC signals (ports 1,2 and ports 3,4) for the receive side. This configuration 5 enables the independent adding and dropping of one of a pair of GbE signals that use the same wavelength. An example is shown in Figure 5. A "west" interface 70 receives an OC-48 signal carrying two GbE signals, which are referred to as GbE-1 and GbE-2. As shown, the GbE-1 10 signal is dropped locally, and the GbE-2 signal is forwarded via the electrical inter-card interface to an "east" interface 72. The east interface 72 adds a new GbE signal, referred to as GbE-3, onto the same wave. The OC-48 signal emerging from the east interface 72 15 includes the signals GbE-2 and GbE-3, as shown. Figure 5 shows the paths for the eastbound signals only. It will be clear that an analogous arrangement can simultaneously be obtained for the westbound signals, so that the westbound GbE-3 is dropped by interface card 72 20 and the westbound GbE-1 is added by the interface card 70. This functionality is achieved through the electrical connection between adjacent cards, as indicated in Figures 3 and 4.

Figures 6 and 7 illustrate a transmitter and 25 receiver, respectively, that use forward error correction (FEC) coding to improve the bit error rate (BER) performance of the GbE streams. A FEC encoder 80 and matching FEC decoder 82, operating at the OC-48 rate, protect the GbE traffic against channel 30 degradations, such as amplifier spontaneous emission (ASE) noise, receiver thermal noise, fiber non-linearities, linear dispersion, etc.

-14-

Methods and apparatus for transporting GbE and FC signals in WDM systems have been described. It will be apparent to those skilled in the art that other modifications to and variations of the above-described 5 technique are possible without departing from the inventive concepts disclosed herein. Accordingly, the invention should be viewed as limited solely by the scope and spirit of the appended claims.

-15-

CLAIMS

What is claimed is:

5 1. A method of transporting a plurality of encoded data signals over a communications link having a data-carrying capacity less than the sum of the respective signaling rates of the encoded data signals, comprising:

10 at a transmitter coupled to the link, (i) decoding the encoded data signals to generate a corresponding plurality of non-encoded data signals, the sum of the respective signaling rates of the non-encoded data signals being no more than the data-carrying capacity of the link, (ii) multiplexing the non-encoded data signals together to form a multiplexed data signal, and (iii) transmitting the multiplexed data signal over the communications link; and

15 at a receiver coupled to the link, (i) receiving the multiplexed data signal transmitted by the transmitter, (ii) de-multiplexing the received signal to obtain the non-encoded data signals, and (iii) re-encoding each de-multiplexed data signal to re-create the encoded data signals.

20 2. A method according to claim 1, wherein the encoded data signals are gigabit Ethernet signals encoded with an 8b/10b encoding.

25 3. A method according to claim 1, wherein the encoded data signals are gigabit Ethernet signals each encoded

-16-

with an 8b/10b encoding, and wherein the multiplexed data signal is an optical carrier signal.

5 4. A method according to claim 3, wherein the number of gigabit Ethernet streams is two, and the optical carrier signal is an optical carrier 48 (OC-48) signal.

10 5. A method according to claim 3, wherein the number of gigabit Ethernet streams is eight, and the optical carrier signal is an optical carrier 192 (OC-192) signal.

15 6. A method according to claim 1, further comprising:
at the transmitter, converting sequential data blocks of each non-encoded signal into corresponding packets, and asynchronously interleaving the packets of the non-encoded signals to create the multiplexed data signal; and

20 20 at the receiver, de-interleaving the packets of the received multiplexed data signal to create corresponding sequential streams of packets, and converting the sequential packets of each stream into the corresponding non-encoded data signal.

25 7. A method according to claim 6, wherein the packets are of a fixed size.

30 8. A method according to claim 1, wherein the encoded data signals are gigabit Ethernet signals encoded with an 8b/10b encoding, and wherein the multiplexed data signal is an optical carrier 48 (OC-48) signal, and further comprising:

-17-

5 at the transmitter, converting sequential data blocks of each non-encoded signal into corresponding fixed-size packets, and asynchronously interleaving the packets of the non-encoded signals to create the OC-48 signal; and

10 at the receiver, de-interleaving the packets of the received multiplexed data signal to create corresponding sequential streams of packets, and converting the sequential packets of each stream into the corresponding non-encoded data signal.

15 9. A method according to claim 1, wherein the encoded data signals are encoded according to a run-length limited code.

10 10. A method according to claim 1, wherein the communications link employs error control coding to detect the occurrence of errors in the multiplexed data signal.

20 11. A method according to claim 1, further comprising:

25 at the transmitter, performing forward error correction (FEC) coding of the multiplexed data signal prior to its transmission over the communications link; and

at the receiver, performing FEC decoding of the received multiplexed data signal.

12. An interface in a network device, comprising:

30 transmitter circuitry operative to (i) receive a plurality of encoded data signals for transmission over a communications link having a data-carrying capacity

-18-

less than the sum of the respective signaling rates of the encoded data signals, (ii) decode the encoded data signals to generate a corresponding plurality of non-encoded data signals, the sum of the respective 5 signaling rates of the non-encoded data signals being no more than the data-carrying capacity of the communications link, (iii) multiplex the non-encoded data signals together to form a multiplexed data signal, and (iv) transmit the multiplexed data signal over the 10 communications link; and

receive circuitry operative to (i) receive a multiplexed data signal from the communications link, the multiplexed data signal including a plurality of data signals multiplexed together, each data signal 15 having been generated by decoding a corresponding encoded data signal, (ii) de-multiplex the received signal to obtain the plurality of data signals, and (iii) encode each de-multiplexed data signal to re-create the encoded data signals.

20

13. An interface according to claim 12, wherein the encoded data signals are gigabit Ethernet signals encoded with an 8b/10b encoding.

25

14. An interface according to claim 12, wherein the encoded data signals are gigabit Ethernet signals encoded with an 8b/10b encoding, and wherein the multiplexed data signal is an optical carrier 48 (OC-48) signal.

30

-19-

15. An interface according to claim 12, wherein:

the transmit circuitry is further operative to convert sequential data blocks of each non-encoded signal into corresponding packets, and asynchronously 5 interleave the packets of the non-encoded signals to create the multiplexed data signal; and

10 the receive circuitry is further operative to de-interleave the packets of the received multiplexed data signal to create corresponding sequential streams of packets, and convert the sequential packets of each 15 stream into the corresponding non-encoded data signal.

16. An interface according to claim 15, wherein the 20 packets are of a fixed size.

15

17. An interface according to claim 12, wherein the 25 encoded data signals are gigabit Ethernet signals encoded with an 8b/10b encoding, and wherein the multiplexed data signal is an optical carrier 48 (OC-48) signal, and wherein:

the transmit circuitry is further operative to convert sequential data blocks of each non-encoded signal into corresponding fixed-size packets, and asynchronously 20 interleave the packets of the non-encoded signals to create the OC-48 signal; and

the receive circuitry is further operative to de-interleave the packets of the received multiplexed data signal to create corresponding sequential streams of packets, and convert the sequential packets of each 30 stream into the corresponding non-encoded data signal.

-20-

18. An interface according to claim 12, wherein the encoded data signals are encoded according to a run-length limited code.

5 19. An interface according to claim 12, wherein the communications link employs error control coding to detect the occurrence of errors in the multiplexed data signal.

10 20. An interface according to claim 12, wherein:
the transmit circuitry is further operative to perform forward error correction (FEC) coding of the multiplexed data signal prior to its transmission over the communications link; and

15 the receive circuitry is further operative to perform FEC decoding of the received multiplexed data signal.

21. An interface according to claim 12 forming a first
20 interface in a pair of interfaces, wherein the plurality of encoded data signals received by the transmit circuitry are a first plurality of received encoded data signals and the plurality of encoded signals re-created by the receive circuitry are a first plurality of
25 re-created encoded data signals, and wherein:

the transmit circuitry is further operative to (i) receive a second plurality of encoded data signals from a second interface in the pair of interfaces, (ii) decode each of the second plurality of encoded data signals to generate a corresponding plurality of non-encoded data signals, and (iii) select either the non-encoded data signals generated from the first

-21-

plurality of encoded data signals or the non-encoded data signals generated from the second plurality of encoded data signals for multiplexing to form the multiplexed data signal; and

5 the receive circuitry is further operative to (i) select either the de-multiplexed data signals or a plurality of data signals received from the second interface for encoding to re-create a second plurality of re-created encoded data signals.

1/7

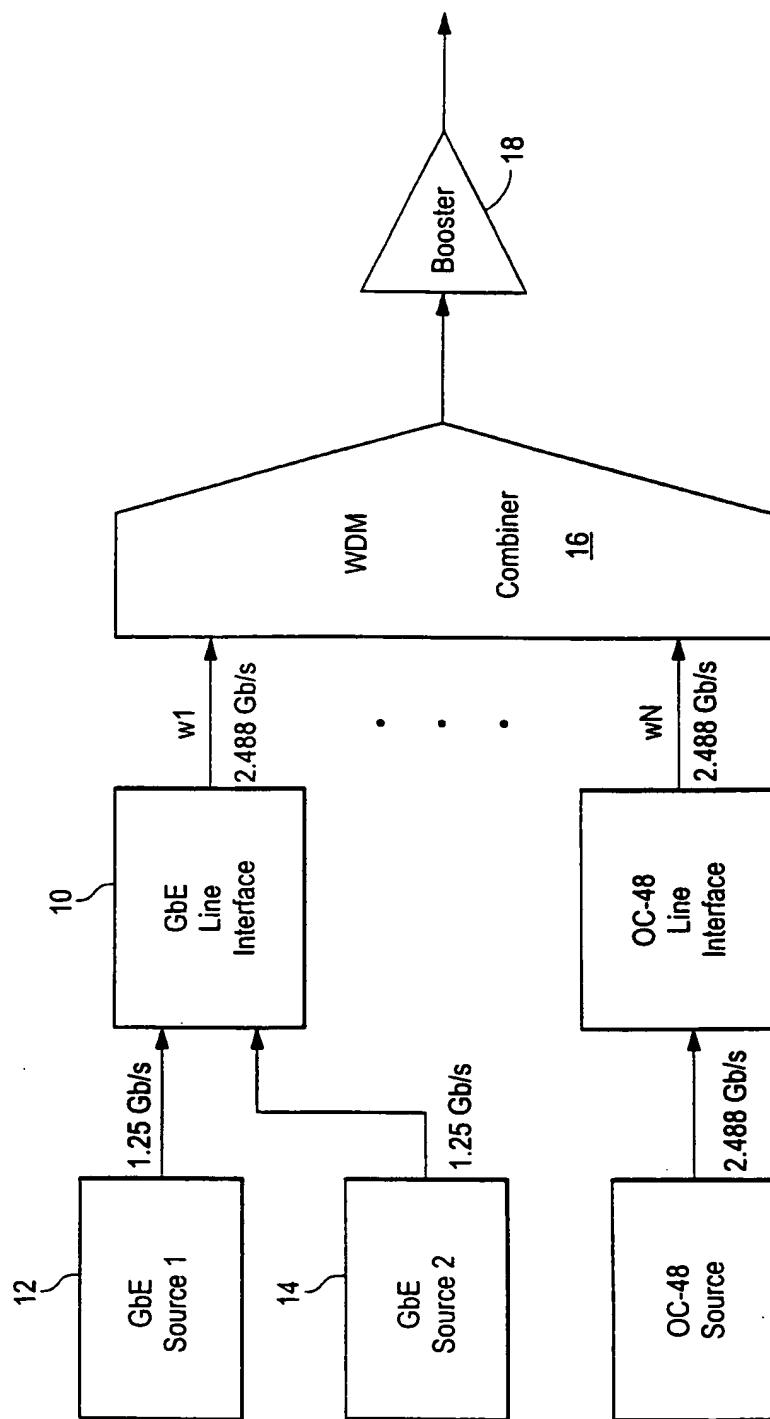
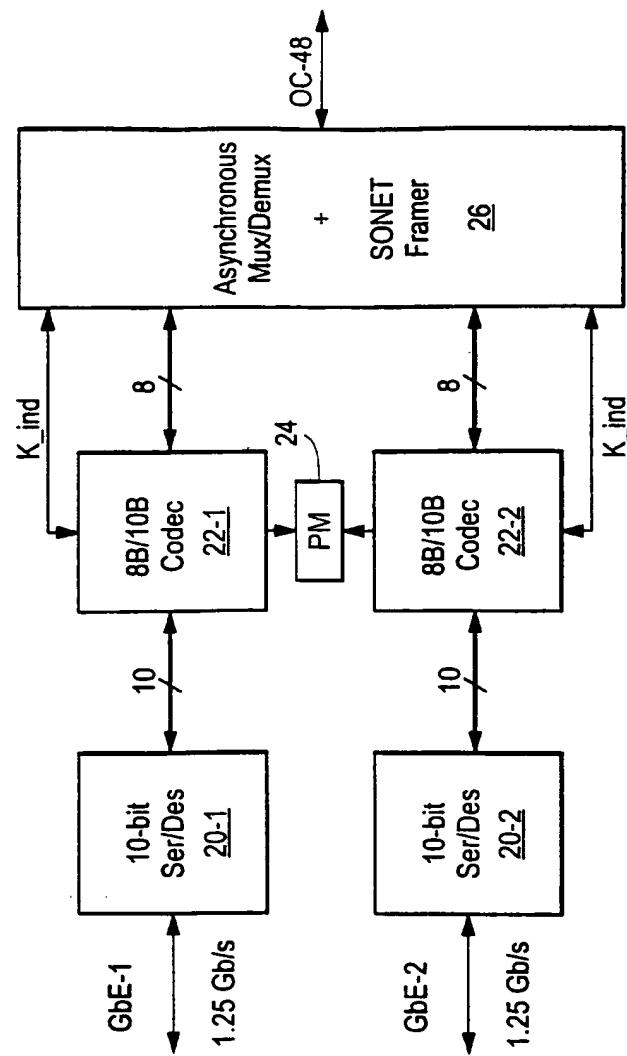


FIG. 1

**FIG. 2**

3/7

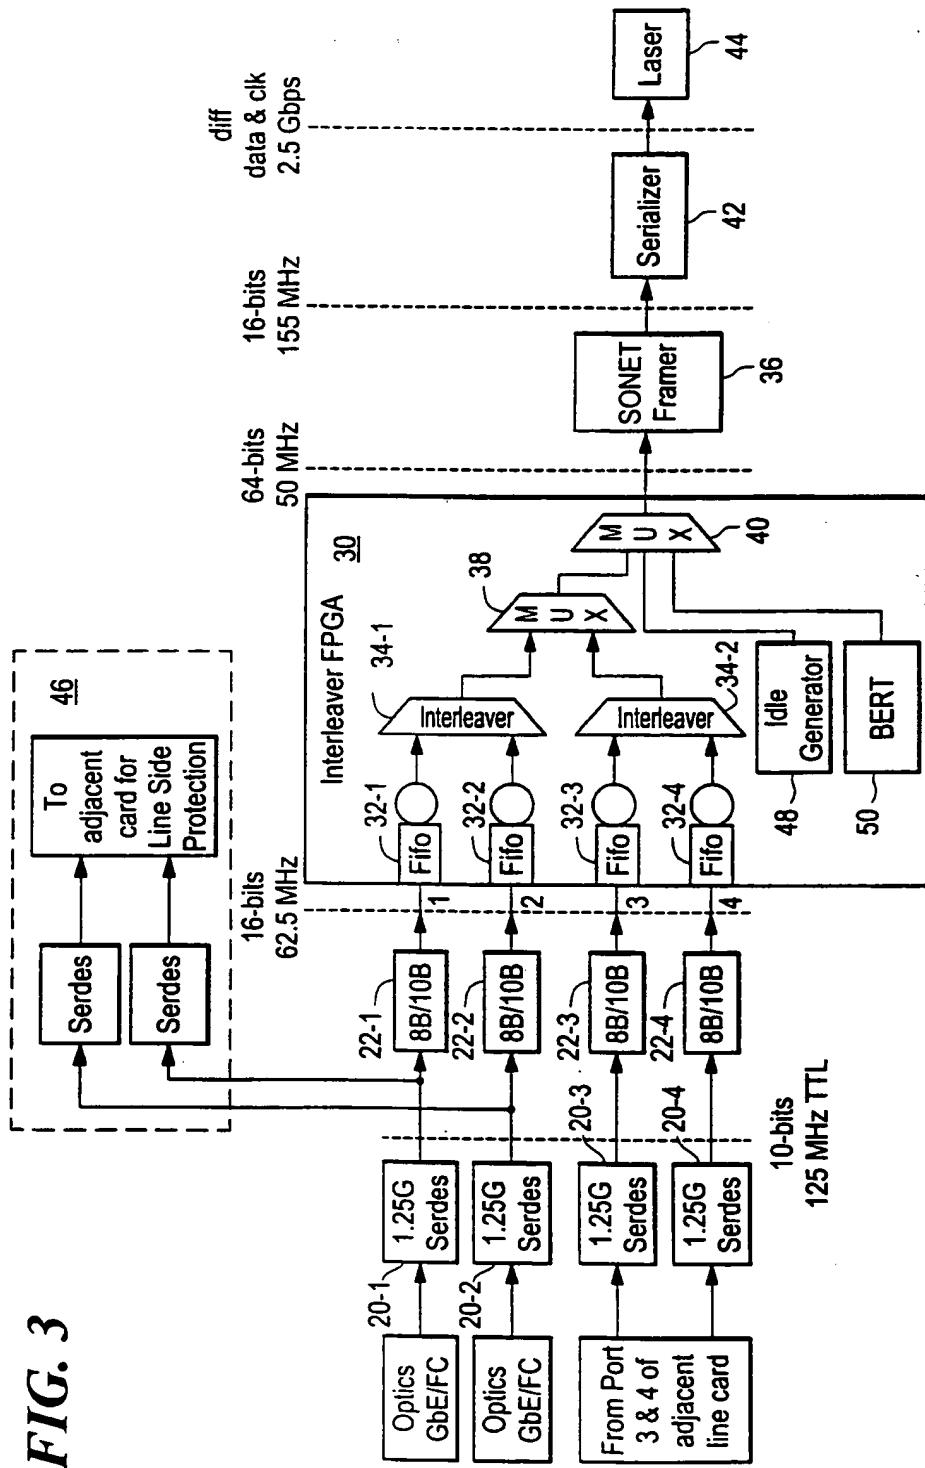
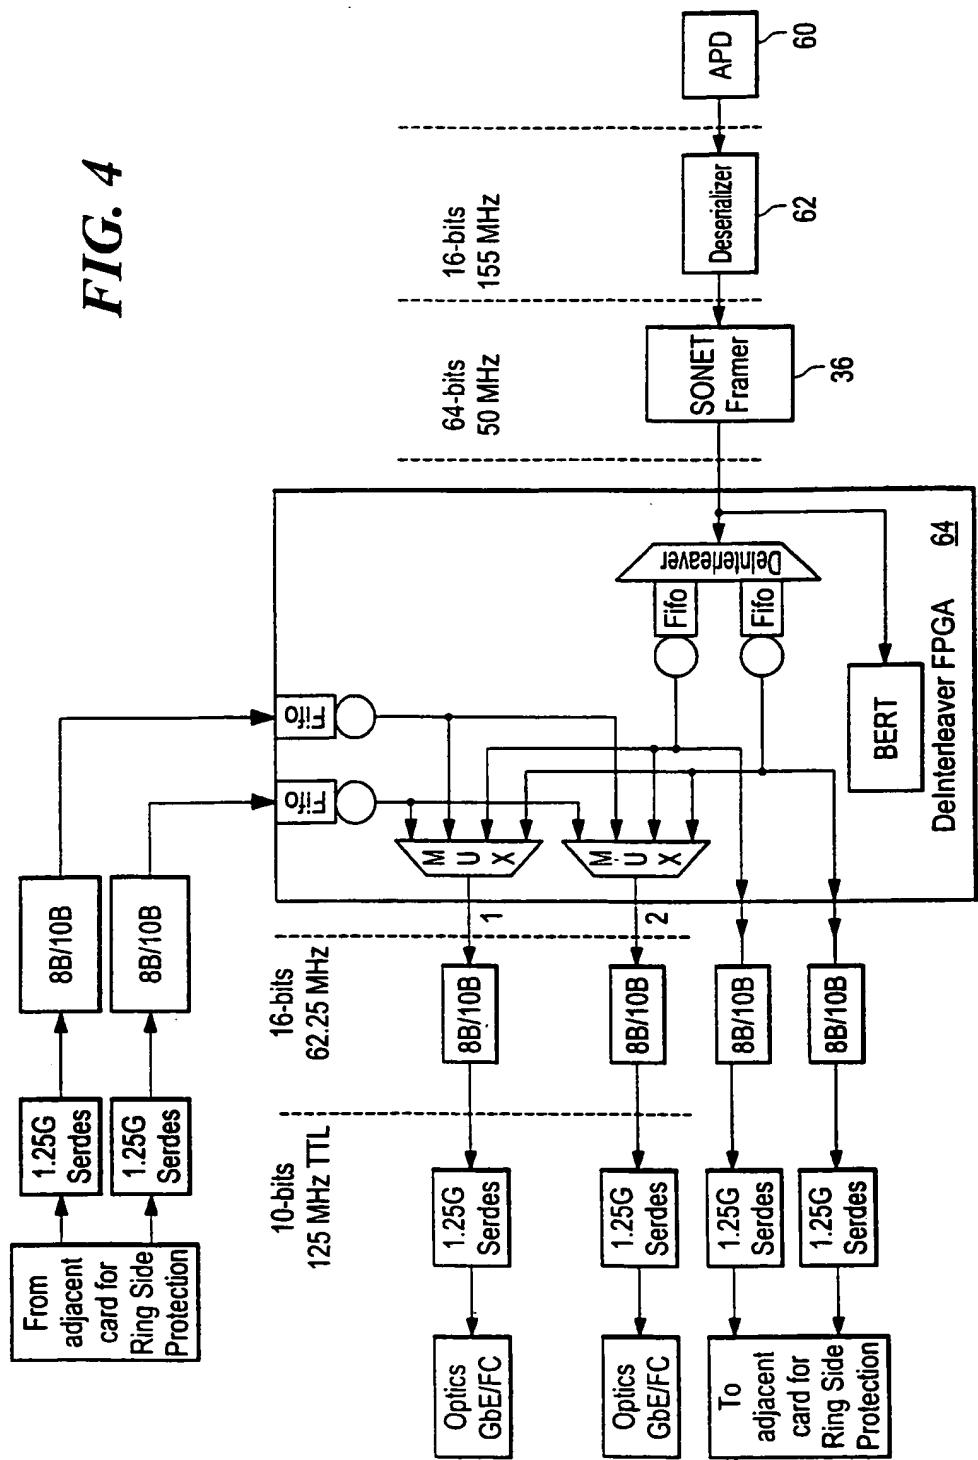


FIG. 4



5/7

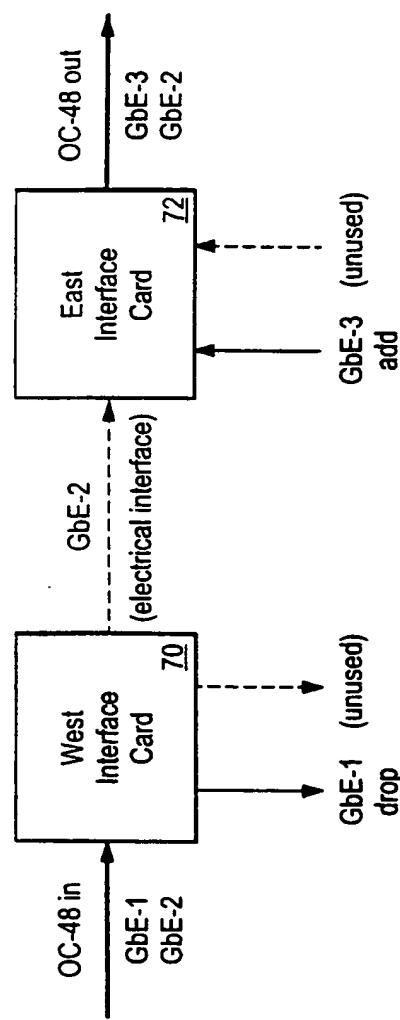


FIG. 5

6/7

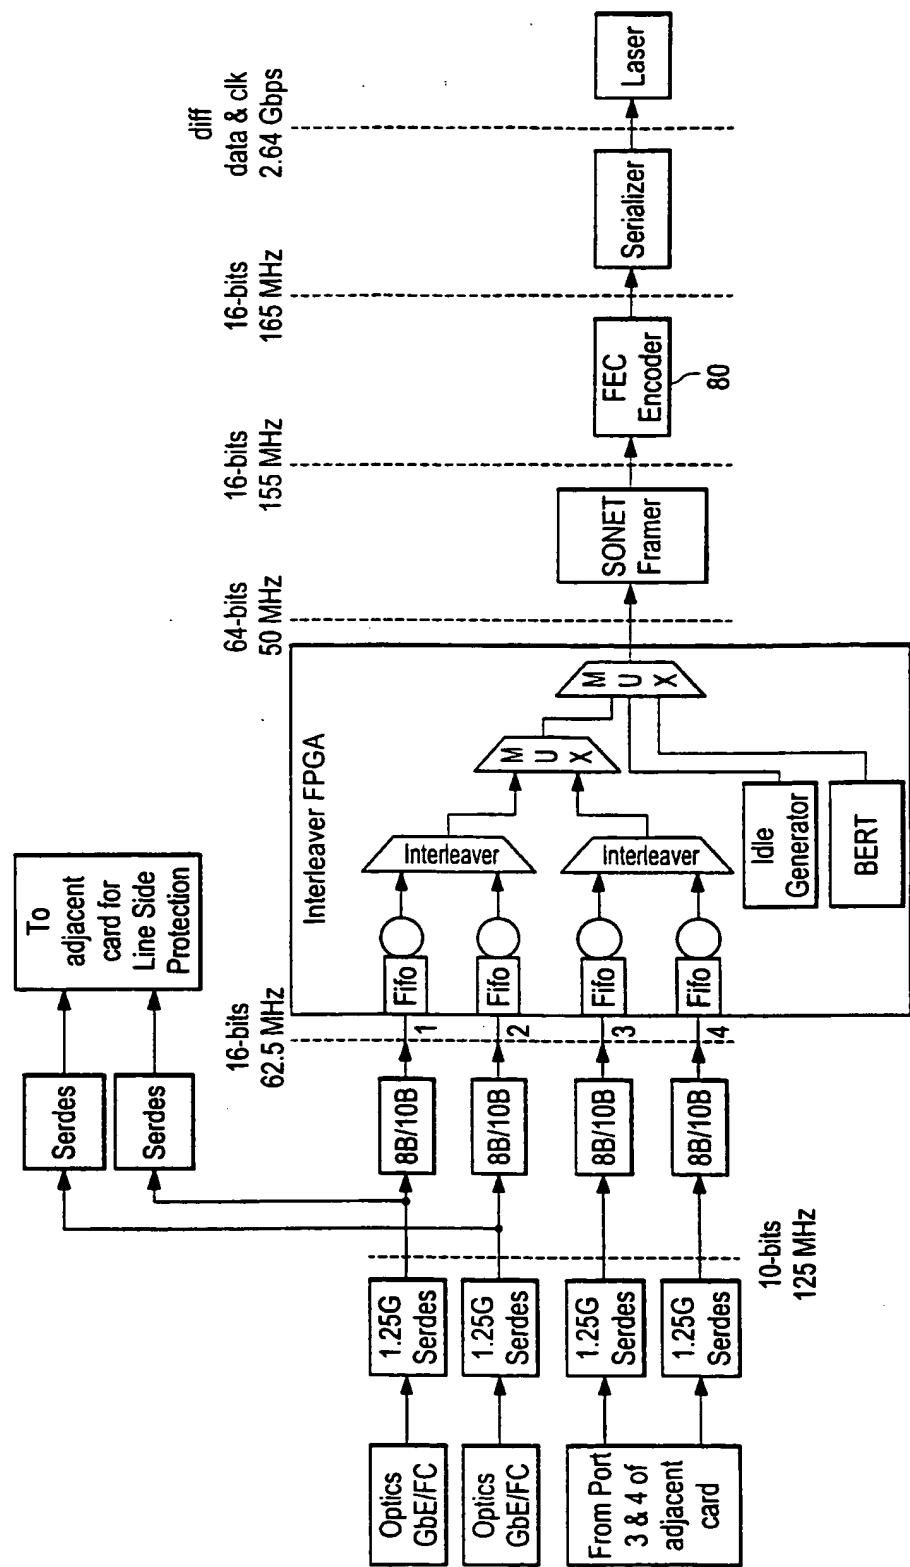
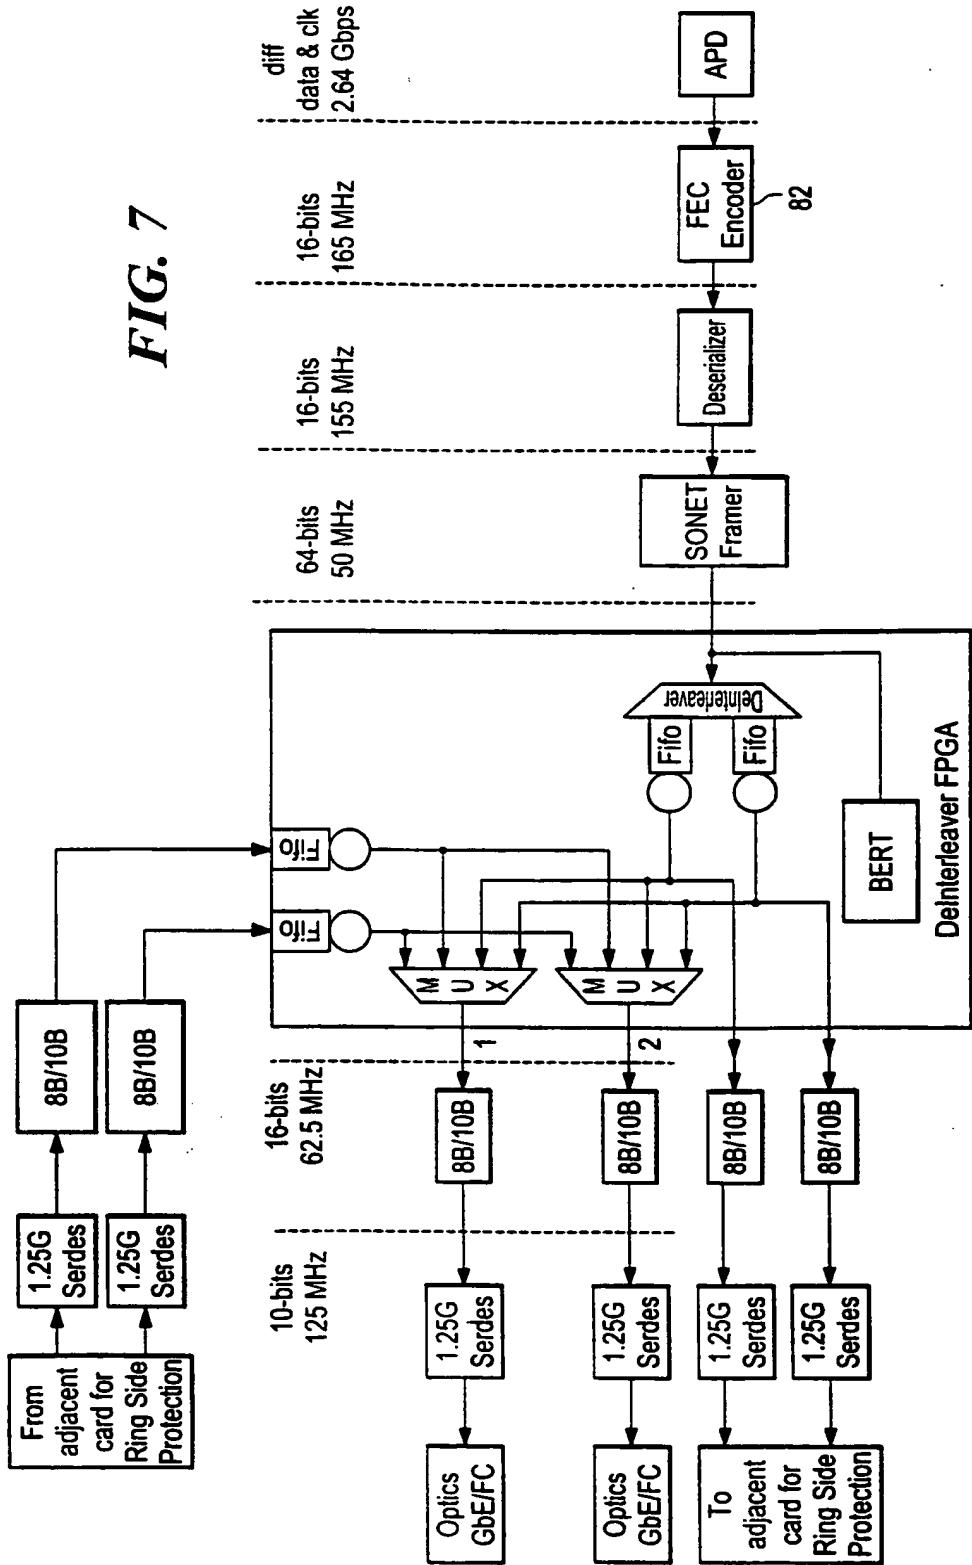


FIG. 6

7/7

FIG. 7



INTERNATIONAL SEARCH REPORT

International application No
PCT/US00/34173

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04J 3/04

US CL : 370/335, 336, 337, 465, 389, 401; 359/135, 126, 136, 118, 115, 114, 117

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/335, 336, 337, 465, 389, 401; 359/135, 126, 136, 118, 115, 114, 117

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

NONE

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

NONE

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,757,416 A (BRICH ET AL.) 26 May 1998, see details of figure 4.	1-21

 Further documents are listed in the continuation of Box C. See patent family annex.

• Special categories of cited documents	*T*	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance		
B earlier document published on or after the international filing date	*X*	document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Y*	document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
O document referring to an oral disclosure, use, exhibition or other means	*A*	document member of the same patent family
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search 20 FEBRUARY 2001	Date of mailing of the international search report 03 APR 2001
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230	Authorized officer DANG TON Telephone No. (703) 305-3230 